

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES**

In Re Application of:

Date: June 20, 2007

Kevin S. BEYER et al.

Confirmation No: 7077

Serial No: 10/788,556

Group Art Unit: 2163

Filed: February 27, 2004

Examiner: Helene Roberta Rose

Title: ASYNCHRONOUS PEER-TO-PEER DATA REPLICATION

Mail Stop Appeal Brief - Patents
Commissioner for Patents
P. O. Box 1450
Alexandria, VA 22313-1450

BRIEF ON APPEAL

(1) Real Party in Interest

The real party in interest is International Business Machines Corporation by virtue of an assignment from the inventors recorded in the U.S. Patent Office on July 21, 2004, reel no. 014888, frame no. 0216.

(2) Related Appeals and Interferences

There are no appeals, interferences, or judicial proceedings known to Appellant, the Appellant's legal representative, or Assignee, which may be related to, directly affect, be directly affected by, or have a bearing on the decision by the Board of Patent Appeals and Interferences in the pending appeal.

(3) Status of Claims

Claims 1-39 and 50-87 stand rejected under 35 U.S.C. § 112, second paragraph, as being indefinite.

Claims 1-5, 7-47, and 50-88 stand rejected under 35 U.S.C. § 102(b) as being anticipated by U.S. Patent No. 5,806,075 (“Jain”).

Claims 1-47 and 50-88 are being appealed.

(4) Status of Amendments

There are no unentered amendments.

(5) Summary of Claimed Subject Matter

In a relational database management system, data is typically stored in tables that have a plurality of rows (or records), in which each row generally has a plurality of columns (or fields). A relational database system can include multiple databases such that one or more copies of a given table residing a particular database (of the relational database system) can be maintained in other databases within the relational database system. In such a relational database system, it is desirable that changes (e.g., inserts, deletes, and updates) to rows in a given table of a database be copied or replicated to other copies of the table that reside in other databases.

Accordingly, independent claim 1 recites a computer-implemented method for providing convergence of multiple copies of a table to a same state in a database system, in which the database system includes a plurality of nodes each having a corresponding copy of the table. The method includes, for each row of each table copy: associating a timestamp with the row, in which the timestamp indicates a time when a change to the row has occurred (specification, page 8, lines 20-21); associating a copy identification to the row, in which the copy identification is an identifier that uniquely identifies the table copy to which the row belongs (specification, page 9, lines 6-7); and associating

propagation controls with the row, in which the propagation controls indicate whether a change to the row should be communicated to other table copies based at least in part on the timestamp of the change or the copy identification associated with the row (specification, page 16, lines 20-21; page 17, lines 10-14).

The method further includes asynchronously capturing a change to a row of a given table copy from a database recovery log, in which the database recovery log contains an entry that describes the change to the row of the given table copy (specification, page 5, line 22 – page 6, line 14; page 17, lines 17-18). The method further includes determining that the captured change to the row of the given table copy is to be communicated to other table copies in the database system, wherein the determination is made in accordance with the indication of the propagation controls associated with the changed row of the given table copy (specification, page 18, lines 2-5). The method further includes communicating the captured change to the other table copies in the database system (specification, page 18, lines 25-26). The method further includes applying the communicated change to the other table copies in the database system, wherein each table copy in the database system converges to a same state (specification, page 19, lines 10-12).

Independent claim 40 recites a database system comprising a plurality of nodes, wherein each node has a corresponding copy of a table (specification, page 5, lines 8-11; FIG. 1). Each row of each table copy includes a timestamp with the row indicating a time when a change to the row has occurred (specification, page 8, lines 20-21), a copy identification, in which the copy identification is an identifier that uniquely identifies the table copy to which the row belongs (specification, page 9, lines 6-7), and propagation

controls that indicate whether a change to the row should be communicated to other table copies based at least in part on the timestamp of the change or the copy identification associated with the row (specification, page 16, lines 20-21; page 17, lines 10-14).

The database system further includes a mechanism to asynchronously capture a change to a row of a given table copy in the database system from a database recovery log, in which the database recovery log contains an entry that describes the change to the row of the given table copy (specification, page 5, line 22 – page 6, line 14; page 17, lines 17-18). The database system further includes a message queue for communicating the captured change to other table copies in the database system (specification, page 5, lines 15-16). The database system further includes a mechanism to apply the communicated change to the other table copies in the database system, wherein each table copy in the database system converges to a same state (specification, page 6, lines 15-16).

Independent claim 50 recites a computer readable medium with program instructions tangibly stored thereon for providing convergence of multiple copies of a table to a same state in a database system, in which the database system includes a plurality of nodes each having a corresponding copy of the table. The computer readable medium comprises instructions for, for each row of each table copy: associating a timestamp with the row, in which the timestamp indicates a time when a change to the row has occurred (specification, page 8, lines 20-21); associating a copy identification to the row, in which the copy identification is an identifier that uniquely identifies the table copy to which the row belongs (specification, page 9, lines 6-7); and associating propagation controls with the row, in which the propagation controls indicate whether a change to the row should be communicated to other table copies based at least in part on

the timestamp of the change or the copy identification associated with the row (specification, page 16, lines 20-21; page 17, lines 10-14).

The computer readable medium further comprises instructions for asynchronously capturing a change to a row of a given table copy from a database recovery log, in which the database recovery log contains an entry that describes the change to the row of the given table copy (specification, page 5, line 22 – page 6, line 14; page 17, lines 17-18). The computer readable medium further comprises instructions for determining that the captured change to the row of the given table copy is to be communicated to other table copies in the database system, wherein the determination is made in accordance with the indication of the propagation controls associated with the changed row of the given table copy (specification, page 18, lines 2-5). The computer readable medium further comprises instructions for communicating the captured change to the other table copies in the database system (specification, page 18, lines 25-26). The computer readable medium further comprises instructions for applying the communicated change to the other table copies in the database system, wherein each table copy in the database system converges to a same state (specification, page 19, lines 10-12).

(6) Grounds of Rejection to be Reviewed on Appeal

1. Appellant requests review as to 1-39 and 50-87 and their rejection under 35 U.S.C. § 112, second paragraph, as being indefinite.
2. Appellant requests review as to claims 1-5, 7-47, and 50-88 and their rejection under 35 U.S.C. § 102(b) as being anticipated by Jain.

(7) **Argument**

1. Claims 1-39 and 50-87 are not properly rejected under 35 U.S.C. § 112, second paragraph, as being indefinite.

In rejecting claims 1-39 and 50-87 under 35 U.S.C. § 112, second paragraph, the Examiner asserts that the limitation “at least in part” is indefinite because neither the claims nor the specification explains what “at least in part” means.

Appellant respectfully disagrees.

First, the limitation “at least in part” is clear on its own terms as set forth in the claim. In particular, claim 1 recites a method that includes (for each row of a table copy) associating a timestamp with the row, associating a copy identification with the row, and associating propagation control with the row. The propagation controls indicate whether a change to the row should be communicated to other table copies based at least in part on the timestamp of the change or the copy identification associated with the row. That is, the propagation control utilizes *the timestamp of the change or the copy identification associated with the row* to indicate whether a change to the row should be communicated to other table copies.

Second, the claims and the specification provide support for such a limitation. More specifically, claim 6 recites that the propagation controls comprise a delete label and a conflict label. As described in the specification on page 17, lines 10-14, the delete label corresponds to a CopyDelete flag 203 (FIG. 2), and the conflict label corresponds to an ImplicitDelete flag 204 (FIG. 2). With reference to the ImplicitDelete flag 204, for example, the ImplicitDelete flag is set when: 1) a row is implicitly deleted and the copy identification of the row is the same as the copy identification assigned to the source table

copy; or 2) when the timestamp of a delete message is higher than a timestamp of an existing row, and the copy identification of the existing row is the same as the copy identification assigned to the target table copy (see specification, page 10 lines 18-23). Accordingly, the ImplicitDelete flag 204 is set *based on the timestamp of a change to the row or the copy identification assigned to the row*. Similarly, the CopyDelete flag 203 is set *based on a copy identification of a row* to prevent propagation of a row delete to other table copies from the treated as a new row delete.

For these reasons, Appellant respectfully requests the Board to reverse the Examiner's rejections of the claims under 35 U.S.C. § 112, second paragraph.

2. Claims 1-5, 7-47, and 50-88 are not properly rejected under 35 U.S.C. § 102(b) as being anticipated by Jain.

Claim 1 recites asynchronously capturing a change to a row of a given table copy from a database recovery log, in which the database recovery log contains an entry that describes the change to the row of the given table copy.

A. Jain Fails To Disclose Asynchronously Capturing A Change To A Row Of A Given Table Copy From A Database Recovery Log As Recited In Claim 1

Jain discloses a system and method for replicating modifications made at a local site to multiple remote sites in a peer-to-peer environment (see Abstract). With respect to data modification, identification, and retention, Jain discloses that in order to propagate the modifications made to data items in one database to the same data items in another database, it is necessary to retain the modifications until they can be made to the other sites (col. 7, ll. 4-7). Jain discloses that prior art methods – e.g., master/slave prior art systems – use a transactional recovery log (i.e., a redo log) to retain and identify database

modifications for propagation (col. 7, ll. 7-9). However, unlike these prior art systems, Jain encodes propagation information directly in tables in the database system. That is, the tables themselves contain all the information necessary to replicate a data modification to other data sites (col. 7, ll. 28-38).

In rejecting claim 1, on page 5 of the Action mailed September 29, 2006, the Examiner cites column 7, lines 7-36 which teaches the redo log. As discussed above, the redo log, however, corresponds to prior art methods in which changes are made only to a master copy, and then transmitted to slave copies. In contrast, Jain discloses techniques for encoding all the information necessary to replicate a data modification within a table of the database system. That is, the tables themselves contain all the information necessary to replicate a data modification to other data sites (col. 7, ll. 28-38).

On page 31 of the Action mailed September 29, 2006, the Examiner asserts that the limitation “asynchronously capturing a change to a row of a given table copy from a database recovery log” is not recited in the rejected claims. However, claim 1 was previously amended to include the limitation and, therefore, such a limitation is not being read from the specification into the claims as asserted by the Examiner.

For the reasons discussed above, Appellant respectfully submits that claim 1 (and the claims that depend therefrom) are, therefore, allowable over Jain.

B. The Examiner has not established anticipation under 35 U.S.C. §102

Anticipation under 35 U.S.C. §102 requires the disclosure in a single piece of prior art of each and every limitation of a claimed invention. *Electro Med. Sys. S.A. v. Cooper Life Sciences*, 34 F.3d 1048, 32 USPQ2d 1017, 1019 (Fed. Cir. 1994). The Examiner has failed to show that Jain discloses asynchronously capturing a change to a

row of a given table copy from a database recovery log, in which the database recovery log contains an entry that describes the change to the row of the given table copy. Thus, claims 1 (and the claims that depend therefrom) are improperly rejected under 35 U.S.C. §102(b) as being anticipated by Jain.

C. Other Independent Claims

Independent claims 40 and 50 each incorporates limitations similar to those of claim 1. Claims 40 and 50, and the claims that depend therefrom, are also allowable over Jain for reasons corresponding to those set forth with respect to claim 1.

Conclusion

Jain fails to disclose asynchronously capturing a change to a row of a given table copy from a database recovery log, in which the database recovery log contains an entry that describes the change to the row of the given table copy. Appellant, therefore, respectfully submits that independent claims 1, 40, and 50 (and the claims that depend therefrom) are not properly rejected under § 102.

Please charge any fee that may be necessary for the continued pendency of this application to Deposit Account No. 09-0460 (IBM Corporation).

Respectfully submitted,
SAWYER LAW GROUP LLP

June 20, 2007

Date

/Kelvin M. Vivian/

Kelvin M. Vivian
Attorney for Appellant
Reg. No. 53,727
(650) 475-1448

Appendix of Claims

1. (Previously Presented) A computer-implemented method for providing convergence of multiple copies of a table to a same state in a database system, the database system including a plurality of nodes each having a corresponding copy of the table, the method comprising:

for each row of each table copy,

associating a timestamp with the row, the timestamp indicating a time when a change to the row has occurred;

associating a copy identification to the row, the copy identification being an identifier that uniquely identifies the table copy to which the row belongs; and

associating propagation controls with the row, the propagation controls indicating whether a change to the row should be communicated to other table copies based at least in part on the timestamp of the change or the copy identification associated with the row;

asynchronously capturing a change to a row of a given table copy from a database recovery log, the database recovery log containing an entry that describes the change to the row of the given table copy;

determining that the captured change to the row of the given table copy is to be communicated to other table copies in the database system, the determination being made in accordance with the indication of the propagation controls associated with the changed row of the given table copy;

communicating the captured change to the other table copies in the database system; and

applying the communicated change to the other table copies in the database system, wherein each table copy in the database system converges to a same state.

2. (Previously Presented) The method of claim 1, wherein the timestamp comprises a monotonic number having a non-decreasing time value, wherein the time values for each table copy is in a common time base.

3. (Previously Presented) The method of claim 2, wherein associating a timestamp with the row includes associating the timestamp with the row in response to a user induced change of the row.

4. (Previously Presented) The method of claim 1, wherein the copy identification uniquely identifying each table copy has an ordering property.

5. (Previously Presented) The method of claim 4, wherein associating a copy identification to the row includes associating the copy identification to the row in response to a user induced change of the row.

6. (Original) The method of claim 1, wherein the propagation controls comprise a delete label and a conflict label.

7. (Previously Presented) The method of claim 6, wherein the delete label indicates that a row delete in a given table copy is not to be communicated to other table copies in the database system.
8. (Previously Presented) The method of claim 6, wherein the conflict label indicates that an implicit row delete in a given table copy is to be communicated to other table copies in the database system, an implicit row delete being a deletion of a row due to a conflict.
9. (Previously Presented) The method of claim 6, wherein asynchronously capturing a change to a row of a given table copy comprises:
 - accessing the entry of the database recovery log pertaining to the change of the row;
 - determining a type of change to the row;
 - extracting old column values and new column values of the row from the entry in the database recovery log; and
 - extracting changed key column values and unchanged key column values associated with the row from the entry in the database recovery log.
10. (Previously Presented) The method of claim 9, wherein the type of change to the row comprises a row insert, a row delete, a non-key update, or a key update.

11. (Previously Presented) The method of claim 10, wherein determining that the captured change to the row of the given table copy is to be communicated to other table copies in the database system comprises:

examining the type of change to the row of the given table copy, the copy identification associated with the row of the given table copy, and the propagation controls associated with the row of the given table copy.

12. (Previously Presented) The method of claim 11, wherein determining that the captured change to the row of the given table copy is to be communicated to other table copies in the database system further comprises:

determining that the captured change to the row of the given table copy is to be communicated to other table copies in the database system responsive to:

the type of change to the row being a row insert; and

the copy identification associated with the row being the same as a copy identification assigned to the given table copy.

13. (Previously Presented) The method of claim 11, wherein determining that the captured change to the row of the given table copy is to be communicated to other table copies in the database system further comprises:

determining that the captured change to the row of the given table copy is to be communicated to other table copies in the database system responsive to:

the type of change to the row being a row delete; and

the delete label associated with the row indicating that the captured change is to be communicated to the other table copies in the database system.

14. (Previously Presented) The method of claim 11, wherein determining that the captured change to the row of the given table copy is to be communicated to other table copies in the database system further comprises:

determining that the captured change to the row of the given table copy is to be communicated to other table copies responsive to:

the type of change to the row being a non-key update or a key update;

neither the delete label nor the conflict label associated with the row indicates that the captured change is not to be communicated to other table copies in the database system; and

a new copy identification associated with the row being the same as the copy identification assigned to the given table copy.

15. (Previously Presented) The method of claim 1, wherein communicating the captured change to the other table copies in the database system comprises:

sending information associated with the captured change to the other table copies in the database system,

wherein responsive to the type of change to the row being a row insert, the information associated with the captured change comprises key column values, non-key column values, a timestamp, and a copy identification associated with the row of the given table copy.

16. (Previously Presented) The method of claim 1, wherein communicating the captured change to the other table copies in the database system comprises:

 sending information associated with the captured change to the other table copies in the database system,

 wherein responsive to the type of change to the row being a row delete, the information associated with the captured change comprises key column values, a timestamp, and a copy identification associated with the row of the given table copy.

17. (Previously Presented) The method of claim 1, wherein communicating the captured change to the other table copies in the database system comprises:

 sending information associated with the captured change to the other table copies in the database system,

 wherein responsive to the type of change being a non-key update, the information associated with the captured change comprises key column values, new non-key column values, an old timestamp, a new timestamp, an old copy identification, and a new copy identification associated with the row of the given table copy.

18. (Previously Presented) The method of claim 1, wherein communicating the captured change to the other table copies in the database system comprises:

 sending information associated with the captured change to the other table copies in the database system,

wherein responsive to the type of change being a key update, the information associated with the captured change comprises old key column values, new key column values, new non-key column values, an old timestamp, a new timestamp, an old copy identification, and a new copy identification associated with the row of the given table copy.

19. (Previously Presented) The method of claim 1, wherein applying the communicated change to the other table copies in the database system comprises:

detecting a conflict for the communicated change with a row of a target table copy, the target table copy being a table copy among the other table copy in the database system at which changes are to be replicated;

determining a priority for the communicated change responsive to the conflict being detected; and

changing the row of the target table copy in accordance with the detected conflict and the determined priority for the communicated change.

20. (Previously Presented) The method of claim 19, wherein detecting a conflict for the communicated change with a row of a target table copy comprises:

responsive to the type of the communicated change being a row delete, the conflict is detected responsive to:

no row of the target table copy with key column values matching key column values associated with the communicated change is identified; or

for a row of the target table copy with key column values matching the key column values associated with the communicated change,

a timestamp associated with the row of the target table copy does not match the timestamp associated with the communicated change; or

a copy identification of the row of the target table copy does not match a copy identification associated with the communicated change.

21. (Previously Presented) The method of claim 19, wherein detecting a conflict for the communicated change with a row of a target table copy further comprises:

responsive to the type of the communicated change being a row insert, the conflict is detected responsive to:

a row of the target table copy with key column values matching key column values associated with the communicated change is identified.

22. (Previously Presented) The method of claim 19, wherein detecting a conflict for the communicated change with a row of a target table copy further comprises:

responsive to the type of the communicated change being a non-key update, the conflict is detected responsive to:

no row of the target table copy with key column values matching key column values associated with the communicated change is identified; or

for a row of the target table copy with key column values matching the key column values associated with the communicated change,

a timestamp associated with the row of the target table copy does not match an old timestamp associated with the communicated change; or

a copy identification of the row of the target table copy does not match an old copy identification associated with the communicated change.

23. (Previously Presented) The method of claim 19, wherein detecting a conflict for the communicated change with a row of a target table copy further comprises:

responsive to the type of the communicated change being a key update, the conflict is detected responsive to:

no row of the target table copy with key column values matching old key column values associated with the communicated change is identified; or

for a row of the target table copy with key column values matching old key column values associated with the communicated change,

a timestamp associated with the row of the target table copy does not match an old timestamp associated with the communication change; or

a copy identification associated with the row of the target table copy does not match an old copy identification associated with the communicated change; or

a row of the target table copy with key column values matching new key column values associated with the communicated change is identified.

24. (Previously Presented) The method of claim 19, wherein applying the communicated change to the other table copies in the database system comprises:

assigning priority to the communicated change responsive to no conflict being detected between the communicated change and the row of the target table copy.

25. (Previously Presented) The method of claim 19, wherein determining a priority for the communicated change comprises:

assigning priority to the communicated change responsive to the type of the communication change ~~is~~ being a row insert; and

a timestamp associated with the communicated change is greater than a timestamp associated with the conflicting row in the target table copy; or

the timestamp associated with the communicated change is equal to the timestamp associated with the conflicting row, and a copy identification associated with the communication change is greater than a copy identification associated with the conflicting row.

26. (Previously Presented) The method of claim 19, wherein determining a priority for the communicated change comprises:

assigning priority to the communicated change responsive to the type of the communicated change being a row delete and responsive to:

no row in the target table copy matches key column values associated with the communicated change; or

a timestamp associated with the communicated change is greater than a timestamp associated with the conflicting row in the target table copy; or

the timestamp associated with the communicated change is the same as the timestamp associated with the conflicting row and a copy identification associated with the communication change is greater than a copy identification associated with the conflicting row.

27. (Previously Presented) The method of claim 19, wherein determining a priority for the communicated change comprises:

assigning priority to the communicated change responsive to the type of the communicated change being a non-key update and responsive to:

no row in the target table copy matches key column values associated with the communicated change; or

a timestamp associated with the communicated change is greater than a timestamp associated with the conflicting row in the target table copy; or

the timestamp associated with the communicated change is the same as the timestamp associated with the conflicting row and a copy identification associated with the communicated change is greater than a copy identification associated with the conflicting row.

28. (Previously Presented) The method of claim 19, wherein determining a priority for the communicated change if the conflict is detected comprises:

assigning priority to the communicated change responsive to the type of the communicated change being a key update and responsive to:

no row in the target table copy matching old key column values associated with the communicated change is identified; and

no row in the target table copy matching new key column values associated with the communicated change is identified; or

a new timestamp associated with the communicated change is greater than a timestamp associated with the conflicting row in the target table copy with key column values matching new key column values associated with the communicated change; or

the new timestamp and copy identification associated with the communicated change matches the timestamp and copy identification associated with the conflicting row in the target table copy with key column values matching new key column values associated with the communicated change, respectively; or

an old timestamp associated with the communicated change is greater than the timestamp associated with the conflicting row in the target table copy with key columns matching old key column values associated with the communicated change; or

the timestamp associated with the communicated change matches the timestamp associated with the conflicting row and an old copy identification associated with the communicated change is greater than the copy identification associated with the conflicting row with key column values matching old key column values associated with the communicated change; and

no row in the target table copy matching new key column values associated with the communicated change is identified; or

the new timestamp associated with the communicated change is greater than the timestamp associated with the conflicting row with key column values matching the new key column values associated with the communicated change; or

the timestamp associated with the communicated change matches the timestamp associated with the conflicting row and the new copy identification associated with the communicated change is greater than the copy identification associated with the conflicting row with key column values matching the new key column values associated with the communicated change.

29. (Previously Presented) The method of claim 19, wherein changing the row of the target table copy comprises:

controlling propagation of the change applied to the target table copy;
insuring convergence of each table copy in the database system to the same state;
and
installing the communicated change into the target table copy.

30. (Previously Presented) The method of claim 29, wherein controlling propagation of the change applied to the target table copy comprises:

responsive to the type of the communicated change being a row insert, setting a copy identification associated with the applied change to the target table copy to a copy identification associated with the communicated change.

31. (Previously Presented) The method of claim 29, wherein controlling propagation of the change applied to the target table copy comprises:

responsive to the type of the communicated change being a row delete, updating a row of the target table copy with key column values matching key column values associated with the communicated change by setting the delete label to indicate not to propagate the applied change.

32. (Previously Presented) The method of claim 29, wherein controlling propagation of the change applied to the target table copy comprises:

responsive to the communicated change being a non-key or key update, setting a copy identification associated with the applied change to the target table copy to the copy identification associated with the communicated change.

33. (Previously Presented) The method of claim 29, wherein insuring convergence of each table copy in the database system to the same state comprises:

insuring propagation of an implicit delete change in the target table copy by setting a conflict label of the propagation controls associated with the conflicting row when the type of the communicated change is a row insert, row delete, or non-key update

assigned priority, and a copy identification associated with the conflicting row is a copy identification assigned to the target table copy.

34. (Previously Presented) The method of claim 29, wherein insuring convergence of each table copy in the database system to the same state comprises:

insuring propagation of a delete change in the target table by setting a conflict label of the propagation controls associated with the conflicting row with key column values matching old key column values associated with the communicated change, when the type of the communicated change is a key update assigned priority, and a copy identification associated with the conflicting row matches a copy identification assigned to the target table copy; and

insuring propagation of a delete change in the target table copy by setting a conflict label of the propagation controls associated with the conflicting row with key column values matching new key column values associated with the communicated change, when the type of the communicated change is a key update assigned priority, and a copy identification associated with the conflicting row is a copy identification assigned to the target table copy.

35. (Previously Presented) The method of claim 29, wherein insuring convergence of each table copy in the database system to the same state comprises:

recording communicated old key column values, an old timestamp, and an old copy identification associated with a conflicting change in the target table copy in a delete

tombstone, when the type of the communicated change is a conflicting delete or a conflicting update with conflicting communicated old timestamp or copy identification.

36. (Previously Presented) The method of claim 29, wherein insuring convergence of each table copy in the database system to the same state comprises:

checking for matching delete and suppressing application of the communication change responsive to a delete tombstone matching the new key columns, new timestamp, and new copy identification associated with the communicated change being found, when the type of the communicated change is an insert change or an update change with a conflicting insert assigned priority.

37. (Previously Presented) The method of claim 1, further comprising:
reporting each conflicting change.

38. (Previously Presented) The method of claim 37, wherein reporting each conflicting change comprises:

reporting conflicting changes of a row delete, a row insert, or a non-key update change only when priority is not assigned to the communicated change and a copy identification associated with a conflicting row in a target table copy is the copy identification assigned to the target table copy.

39. (Previously Presented) The method of claim 37, wherein reporting each conflicting change comprises:

reporting conflicting changes of a key update change only when priority is not assigned to the communicated change; and

a copy identification associated with a conflicting row in a target table copy with key column values matching old key column values associated with the communicated change is the copy identification assigned to the target table copy; or

a copy identification associated with a conflicting row with key column values matching new key column values associated with the communicated change is the copy identification assigned to the target table copy.

40. (Previously Presented) A database system, comprising:

a plurality of nodes, each node having a corresponding copy of a table, wherein each row of each table copy includes,

a timestamp with the row indicating a time when a change to the row has occurred;

a copy identification, the copy identification being an identifier that uniquely identifies the table copy to which the row belongs; and

propagation controls indicating whether a change to the row should be communicated to other table copies based at least in part on the timestamp of the change or the copy identification associated with the row;

a mechanism to asynchronously capture a change to a row of a given table copy in the database system from a database recovery log, the database recovery log containing an entry that describes the change to the row of the given table copy;

a message queue for communicating the captured change to other table copies in the database system; and

a mechanism to apply the communicated change to the other table copies in the database system, wherein each table copy in the database system converges to a same state.

41. (Previously Presented) The database system of claim 40, wherein the timestamp comprises a monotonic number having a non-decreasing time value, wherein the time values for each table copy is in a common time base.

42. (Previously Presented) The database system of claim 41, wherein the timestamp is associated with the row in response to a user induced change of the row.

43. (Previously Presented) The database system of claim 40, wherein the copy identification assigned to each table copy has an ordering property.

44. (Previously Presented) The database system of claim 40, wherein the copy identification of a given row is associated with the row in response to a user induced change of the row.

45. (Previously Presented) The database system of claim 40, wherein the propagation controls comprise a delete label and a conflict label.

46. (Previously Presented) The database system of claim 45, wherein the delete label indicates that a row delete in any of the plurality of table copies is not to be communicated to other table copies in the database system.

47. (Previously Presented) The database system of claim 45, wherein the conflict label indicates that an implicit row delete in a given table copy is to be communicated to other table copies in the database system, an implicit row delete being a deletion of a row due to a conflict.

48-49. (Cancelled)

50. (Previously Presented) A computer readable medium with program instructions tangibly stored thereon for providing convergence of multiple copies of a table to a same state in a database system, the database system including a plurality of nodes each having a corresponding copy of the table, the computer readable medium comprising instructions for:

for each row of each table copy,

associating a timestamp with the row, the timestamp indicating a time when a change to the row has occurred;

associating a copy identification to the row, the copy identification being an identifier that uniquely identifies the table copy to which the row belongs; and

associating propagation controls with the row, the propagation controls indicating whether a change to the row should be communicated to other table

copies based at least in part on the timestamp of the change or the copy identification associated with the row;

asynchronously capturing a change to a row of a given table copy from a database recovery log, the database recovery log containing an entry that describes the change to the row of the given table copy;

determining that the captured change to the row of the given table copy is to be communicated to other table copies in the database system, the determination being made in accordance with the indication of the propagation controls associated with the changed row of the given table copy;

communicating the captured change to the other table copies in the database system; and

applying the communicated change to the other table copies in the database system, wherein each table copy in the database system converges to a same state.

51. (Previously Presented) The computer readable medium of claim 50, wherein timestamp comprises a monotonic number having a non-decreasing time value, wherein the time values for each table copy is in a common time base.

52. (Previously Presented) The computer readable medium of claim 51, wherein the instructions for associating a timestamp with the row include instructions for associating the timestamp with the row in response to a user induced change of the row.

53. (Previously Presented) The computer readable medium of claim 50, wherein the copy identification uniquely identifying each table copy has an ordering property.

54. (Previously Presented) The computer readable medium of claim 53, wherein the instructions for associating a copy identification to the row include instructions for associating the copy identification to the row in response to a user induced change of the row.

55. (Previously Presented) The computer readable medium of claim 50, wherein the propagation controls comprise a delete label and a conflict label.

56. (Previously Presented) The computer readable medium of claim 55, wherein the delete label indicates that a row delete in a given table copy is not to be communicated to other table copies in the database system.

57. (Previously Presented) The computer readable medium of claim 55, wherein the conflict label indicates that an implicit row delete in a given table copy is to be communicated to other table copies in the database system, an implicit row delete being a deletion of a row due to a conflict.

58. (Previously Presented) The computer readable medium of claim 50, wherein the instructions for asynchronously capturing a change to a row of a given table copy include instructions for:

accessing the entry of the database recovery log pertaining to the change of the row;

determining a type of change to the row;

extracting old column values and new column values of the row from the entry in the database recovery log; and

extracting changed key column values and unchanged key column values associated with the row from the entry in the database recovery log.

59. (Previously Presented) The computer readable medium of claim 58, wherein the type of change to the row comprises a row insert, a row delete, a non-key update, or a key update.

60. (Previously Presented) The computer readable medium of claim 59, wherein the instructions for determining that the captured change to the row of the given table copy is to be communicated to other table copies in the database system include instructions for:

examining the type of change to the row of the given table copy, the copy identification associated with the row of the given table copy, and the propagation controls associated with the row of the given table copy.

61. (Previously Presented) The computer readable medium of claim 60, wherein the instructions for determining that the captured change to the row of the given table copy is to be communicated to other table copies in the database system further include instructions for:

determining that the captured change to the row of the given table copy is to be communicated to other table copies in the database system responsive to:

the type of change to the row being is a row insert; and

the copy identification associated with the row being the same as a copy identification assigned to the given table copy.

62. (Previously Presented) The computer readable medium of claim 60, wherein the instructions for determining that the captured change to the row of the given table copy is to be communicated to other table copies in the database system further include instructions for:

determining that the captured change to the row of the given table copy is to be communicated to other table copies in the database system responsive to:

the type of change to the row being a row delete; and

the delete label associated with the row indicating that the captured change is to be communicated to the other table copies in the database system.

63. (Previously Presented) The computer readable medium of claim 60, wherein the instructions for determining that the captured change to the row of the given table copy is to be communicated to other table copies in the database system further include instructions for:

determining that the captured change to the row of the given table copy is to be communicated to other table copies in the database system responsive to:

the type of change to the row being a non-key update or a key update; and

neither the delete label nor the conflict label associated with the row indicates that the captured change is not to be communicated to other table copies in the database system; and

a new copy identification associated with the row is the same as the copy identification assigned to the given table copy.

64. (Previously Presented) The computer readable medium of claim 50, wherein the instructions for communicating the captured change to the other table copies in the database system include instructions for:

sending information for the captured change to the others of the plurality of table copies,

wherein responsive to the captured change being a row insert, the information comprises key column values, non-key column values, a monotonic number, and a copy identification of the captured change.

65. (Previously Presented) The computer readable medium of claim 50, wherein the instructions for communicating the captured change to the other table copies in the database system include instructions for:

sending information for the captured change to the others of the plurality of table copies,

wherein responsive to the captured change being a row delete, the information comprises key column values, a monotonic number, and a copy identification of the captured change.

66. (Previously Presented) The computer readable medium of claim 50, wherein the instructions for communicating the captured change to the other table copies in the database system include instructions for:

sending information for the captured change to the others of the plurality of table copies,

wherein responsive to the captured change being a non-key update, the information comprises key column values, new non-key column values, an old monotonic number, a new monotonic number, an old copy identification, and a new copy identification of the captured change.

67. (Previously Presented) The computer readable medium of claim 50, wherein the instructions for communicating the captured change to the other table copies in the database system include instructions for:

sending information ~~for~~ associated with the captured change to other table copies in the database system,

wherein responsive to the type of change to the row being a key update, the information associated with the captured change comprises old key column values, new key column values, new non-key column values, an old timestamp, a new timestamp, an old copy identification, and a new copy identification associated with the row of the given table copy.

68. (Previously Presented) The computer readable medium of claim 50, wherein the instructions for applying the communicated change to the other table copies in the database system include instructions for:

detecting a conflict for the communicated change with a row of a target table copy, the target table copy being a table copy among the other table copy in the database system at which changes are to be replicated;

determining a priority for the communicated change responsive to the conflict being detected; and

changing the row of the target table copy in accordance with the detected conflict and the determined priority for the communicated change.

69. (Previously Presented) The computer readable medium of claim 68, wherein the instructions for detecting a conflict for the communicated change with a row of a target table copy include instructions for:

responsive to the type of the communicated change being a row delete, the conflict is detected responsive to:

no row of the target table copy with key column values matching key column values associated with the communicated change is identified; or

for a row of the target table copy with key column values matching the key column values associated with the communicated change,

a timestamp associated with the row of the target table copy does not match the timestamp associated with the communicated change; or

a copy identification of the row of the target table copy does not match a copy identification associated with the communicated change.

70. (Previously Presented) The computer readable medium of claim 68, wherein the instructions for detecting a conflict for the communicated change with a row of a target table copy further include instructions for:

responsive to the type of the communicated change being a row insert, the conflict is detected responsive to a row of the target table copy with key column values matching key column values associated with the communicated change being identified.

71. (Previously Presented) The computer readable medium of claim 68, wherein the instructions for detecting a conflict for the communicated change with a row of a target table copy further include instructions for:

responsive to the type of the communicated change being a non-key update, the conflict is detected responsive to:

no rows of the target table copy with key column values matching key column values associated with the communicated change is identified; or

for a row of the target table copy with key column values matching the key column values associated with the communicated change,

a timestamp associated with the row of the target table copy does not match an old timestamp associated with the communicated change; or

a copy identification of the row of the target table copy does not match an old copy identification associated with the communicated charge.

72. (Previously Presented) The computer readable medium of claim 68, wherein the instructions for detecting a conflict for the communicated change with a row of a target table copy further include instructions for:

responsive to the type of the communicated change being a key update, the conflict is detected responsive to:

no row of the target table copy with key column values matching old key column values associated with the communicated change is identified; or

for a row of the target table copy with key column values matching old key column values associated with the communicated change,

a timestamp associated with the row of the target table copy does not match an old timestamp associated with the communication change; or

a copy identification associated with the row of the target table copy does not match an old copy identification associated with the communicated change; or

a row of the target table copy with key column values matching new key column values associated with the communicated change is identified.

73. (Previously Presented) The computer readable medium of claim 68, wherein the instructions for applying the communicated change to the other table copies in the database system include instructions for:

assigning priority to the communicated change responsive to no conflict being detected between the communicated change and the row of the target table copy.

74. (Previously Presented) The computer readable medium of claim 68, wherein the instructions for determining a priority for the communicated change include instructions for:

assigning priority to the communicated change responsive to the type of the communication change being a row insert; and

a timestamp associated with the communicated change is greater than a timestamp associated with the conflicting row in the target table copy; or

the timestamp associated with the communicated change is equal to the timestamp associated with the conflicting row; and

a copy identification associated with the communication change is greater than a copy identification associated with the conflicting row.

75. (Previously Presented) The computer readable medium of claim 68, wherein the instructions for determining a priority for the communicated change include instructions for:

assigning priority to the communicated change responsive to the type of the communicated change being a row delete and responsive to:

no row in the target table copy matches key column values associated with the communicated change; or

a monotonic timestamp associated with the communicated change is greater than a timestamp associated with the conflicting row in the target table copy; or

the timestamp associated with the communicated change is the same as the timestamp associated with the conflicting row and a copy identification associated with the communication change is greater than a copy identification associated with the conflicting row.

76. (Previously Presented) The computer readable medium of claim 68, wherein the instructions for determining a priority for the communicated change include instructions for:

assigning priority to the communicated change responsive to the type of the communicated change being a non-key update and responsive to:

no row in the target table copy matches key column values associated with the communicated change; or

a timestamp associated with the communicated change is greater than a timestamp associated with the conflicting row in the target table copy; or

the monotonic timestamp associated with the communicated change is the same as the timestamp associated with the conflicting row and a copy identification associated with the communicated change is greater than a copy identification associated with the conflicting row.

77. (Previously Presented) The computer readable medium of claim 68, wherein the instructions for determining a priority for the communicated change include instructions for:

assigning priority to the communicated change responsive to the type of the communicated change being a key update and responsive to:

no row in the target table copy matching old key column values associated with the communicated change is identified; and

no row in the target table copy matching new key column values associated with the communicated change is identified; or

a new timestamp associated with the communicated change is greater than a timestamp associated with the conflicting row in the target table copy with key column values matching new key column values associated with the communicated change; or

the new timestamp and a copy identification associated with the communicated change matches the timestamp and copy identification associated with the conflicting row in the target table copy with key column values matching new key column values associated with the communicated change, respectively; or

an old timestamp associated with the communicated change is greater than the timestamp associated with the conflicting row in the target table copy with key columns matching old key column values associated with the communicated change; or

the timestamp associated with the communicated change matches the timestamp associated with the conflicting row and an old copy identification associated with the communicated change is greater than the copy identification

associated with the conflicting row with key column values matching old key column values associated with the communicated change; and

no row in the target table copy matching new key column values associated with the communicated change is identified; or

the new timestamp associated with the communicated change is greater than the timestamp associated with the conflicting row with key column values matching the new key column values associated with the communicated change; or

the timestamp associated with the communicated change matches the timestamp associated with the conflicting row and the new copy identification associated with the communicated change is greater than the copy identification associated with the conflicting row with key column values matching the new key column values associated with the communicated change.

78. (Previously Presented) The computer readable medium of claim 68, wherein the instructions for changing the row of the target table copy include instructions for:

controlling propagation of the change applied to the target table copy;

insuring convergence of each table copy in the database system to the same state;

and

installing the communicated change into the target table copy.

79. (Previously Presented) The computer readable medium of claim 78, wherein the instructions for controlling propagation of the change applied to the target table copy include instructions for:

responsive to the type of the communicated change being a row insert, setting a copy identification associated with the applied change to the target table copy to a copy identification associated with the communicated change.

80. (Previously Presented) The computer readable medium of claim 78, wherein the instructions for controlling propagation of the change applied to the target table copy include instructions for:

responsive to the type of the communicated change being a row delete, updating a row of the target table copy with key column values matching key column values associated with the communicated change by setting the delete label to indicate not to propagate the applied change.

81. (Previously Presented) The computer readable medium of claim 78, wherein the instructions for controlling propagation of the change applied to the target table copy include instructions for:

responsive to the type of the communicated change being a non-key or key update, setting a copy identification of the applied change to the target table copy to the copy identification received with the communicated change.

82. (Previously Presented) The computer readable medium of claim 78, wherein the instructions for insuring convergence of each table copy in the database system to the same state include instructions for:

insuring propagation of an implicit delete change in the target table copy by setting a conflict label of the propagation controls associated with the conflicting row when the type of the communicated change is a row insert, a row delete, or a non-key update assigned priority and a copy identification associated with the conflicting row is a copy identification assigned to the target table copy.

83. (Previously Presented) The computer readable medium of claim 78, wherein the instructions for insuring convergence of each table copy in the database system to the same state include instructions for:

insuring propagation of a delete change in the target table by setting a conflict label of the propagation controls associated with the conflicting row with key column values matching old key column values associated with the communicated change, when the type of the communicated change is a key update assigned priority and a copy identification associated with the conflicting row matches a copy identification assigned to the target table copy.

84. (Previously Presented) The computer readable medium of claim 78, wherein the instructions for insuring convergence of each table copy in the database system to the same state include instructions for:

recording communicated old key column values, an old timestamp, and an old copy identification associated with a conflicting change in the target table copy in a delete tombstone, when the type of the communicated change is a conflicting delete or a conflicting update with conflicting communicated old timestamp or copy identification.

85. (Previously Presented) The computer readable medium of claim 78, wherein the instructions for insuring convergence of each table copy in the database system to the same state include instructions for:

checking for matching delete and suppressing application of the communication change responsive to a delete tombstone matching new key column values, new timestamp, and new copy identification associated with the communicated change is found, when the type of the communicated change is an insert change or an update change with a conflicting insert assigned priority.

86. (Previously Presented) The computer readable medium of claim 50, further comprising instructions for:

reporting each conflicting change.

87. (Previously Presented) The computer readable medium of claim 86, wherein the instructions for reporting each conflicting change include instructions for:

reporting conflicting changes of a row delete, a row insert, or a non-key update change only when priority is not assigned to the communicated change and a copy

identification ~~for~~ associated with a conflicting row in a target table copy is the copy identification assigned to the target table copy.

88. (Previously Presented) The computer readable medium of claim 86, wherein the instructions for reporting each conflicting change include instructions for:

reporting conflicting changes of a key update change only when priority is not assigned to the communicated change; and

a copy identification associated with a conflicting row in a target table copy with key column values matching old key column values associated with the communicated change is the copy identification assigned to the target table copy; or

a copy identification associated with a conflicting row with key column values matching new key column values associated with the communicated change is the copy identification assigned to the target table copy.

89-90. (Cancelled)

EVIDENCE APPENDIX

None

RELATED PROCEEDINGS APPENDIX

None